AUTHOR QUERY FORM

ELSEVIER	**Journal:** JAL **Article Number:** 233	**Please e-mail or fax your responses and any corrections to:** **E-mail: corrections.essd@elsevier.vtex.lt** **Fax: +1 61 9699 6735**

Dear Author,

Please check your proof carefully and mark all corrections at the appropriate place in the proof (e.g., by using on-screen annotation in the PDF file) or compile them in a separate list. To ensure fast publication of your paper please return your corrections within 48 hours.

For correction or revision of any artwork, please consult *http://www.elsevier.com/artworkinstructions*

Any queries or remarks that have arisen during the processing of your manuscript are listed below and highlighted by flags in the proof. Click on the 'Q' link to go to the location in the proof.

Location in article	Query / Remark: click on the Q link to go Please insert your reply or correction at the corresponding line in the proof
Q1	Please supply keywords (p. 1/ line 23)
Q2, Q3, Q4, Q5, Q6	Refs. [15,19,23,25,27] were noted in the reference list but are missing from the text - please position each reference in the text or delete it from the list (p. 11/ line 60; p. 12/ line 3,6,8,10)

Thank you for your assistance.

Contents lists available at ScienceDirect

Journal of Applied Logic

www.elsevier.com/locate/jal

Prolegomenon to norms in branching space–times ☆

Nuel Belnap

University of Pittsburgh, 1028-A Cathedral of Learning, Pittsburgh, PA, United States

ARTICLE INFO

Article history:
Available online xxxx

Keywords:
???

ABSTRACT

Earlier (collaborative) work postulated that agency is essentially indeterministic, and finds a home in "branching times" (BT). BT is built on world-wide but momentary events called "moments" structured by an indeterministic causal ordering. The notion of a "history" is defined as a suitable chain of moments; histories branch one from another at moments. Other earlier work, the essential points of which are summarized in Section 2, led to "branching space–times" (BST). BT moments are replaced in BST by local "point events," and "histories" are redefined more realistically as a family of suitable space–time-like structures that branch from one another at point events. In BST one finds that "transitions" from initial events to outcome events—each properly defined—play a fundamental role as originating causes (causae causantes) and effects. This spatio-temporal-causal story provides, in Section 3, an objective foundation for an indeterministic theory of agents and choices amid BST (BSTAC), the key concepts of which are "life event," "life history," and "person," the latter essentially involving alternate possibilities, some of which arise from free choices. Section 4 works out a concept of "choice point," leading to a BST account of "agent α sees to outcome O," which is more sophisticated than the *stit* concept of earlier work. Many-agent cooperation is considered in Section 5, with hints concerning message sending and receiving. A theory of norms in BST (BSTACN) that builds on the foregoing, and also on earlier work, is briefly adumbrated in Section 6.

1. Norms and agency

The idea of norms presupposes agency, and agency presupposes an indeterministic causal order. This has been, of course, a far from universal opinion since philosophers such as Kant and Hume, each in his own way a strict deterministic, populated the world with so-called "compatibilists," meaning those who believe that agency is compatible with (or even requires) determinism. Naturally this dialectical situation has given rise to "incompatibilism," where what is at stake is compatibility/incompatibility with determinism. It is no surprise that arguments one way or the other generally turn out fruitless. A good way of avoiding this particular knot is to shift focus to the question of the compatibility between agency and *in*determinism, or even more interestingly, to plunge ahead under the belief that agency *presupposes* indeterminism in the causal order.

This stance, which lines up with that of Kane [14], is taken by Belnap et al. [9] (*Facing the future*, henceforth FF), which concentrates on agents in so-called "branching time" (BT). FF discusses at length how agency fits into and is illuminated by BT, arguing that the rudiments of indeterminism as required by agency are well-modeled by BT. Branching time is not, however, a comfortable background for stating the doctrine that an agentive choice is a locatable event that influences only

☆ Thanks to Thomas Müller for many corrections and comments. His is the eye of an eagle. Two helpful referees each made suggestions that I have been pleased to use.

E-mail address: belnap@pitt.edu.

1570-8683/$ – see front matter © 2011 Published by Elsevier B.V.
doi:10.1016/j.jal.2009.09.006

its own causal future. The rewording of this doctrine as "no action at a distance" makes it clear how alien it is to branching time, which has no concept of distance, nor even the much more modest concept of "over there." That doctrine can only be stated and understood in the context of a causal order that, unlike branching time, supports the idea of "over there," which the jargon of Special Relativity calls "space-like relatedness." A causal order admitting both indeterminism and space-like relatedness is well-modeled, up to an approximation, in "branching space–times," henceforth BST, as defined in Belnap [1].[1] In the normal case, space-like relatedness of choice events guarantees that the choices are independent (not correlated). The abnormal, even weird, case of space-like related choices that are nevertheless correlated is called "funny business"; it is investigated in Belnap [3,5], and Müller et al. [22]. How causes, in the guise of *causae causantes*, or "originating causes," fit into BST, provided there is no funny business, is developed in Belnap [7]. We may take over a good deal of that theory if we identify free choices as among the originating causes in BST. From now on, I'll drop "free" as redundant, in the belief that so-called "Frankfurt cases" equivocate: Given that we suppress vagueness as a red herring, something called a choice is either free, or it's not really a choice. The identification of choice and free choice ensures the *principled* propriety of claiming that certain choices are entirely independent. To this end, we'll outline the beginnings of a theory of "branching space–times with agents and choices" (BSTAC), using previous work as much as possible. We aim for a mathematically rigorous theory; as a consequence, the theory necessarily avoids the typical waffling about agency with mental concepts, and as such does not pretend to offer a "complete" story.

2. Rudiments of the theory of branching space–times

One way of understanding the aim of BST is by thinking of it as arising by two independent moves from Newtonian theory, a theory that is both deterministic and non-relativistic. Move 1: "Branching time" theory comes from the Newtonian basis by replacing determinism with indeterminism, while still being non-relativistic. The "world" of branching time is composed of "moments" in a tree-like arrangement, with each complete branch counting as a possible "history." It is critical that each moment should be pictured as a "simultaneity slice" extending straight across the entire universe, "edge to edge," so to speak. The picture of a history as a linear order of spatially wide instantaneous events violates the relativistic principle that there is no absolute meaning to simultaneity. Furthermore, in branching time, a branch point between two histories is the entire "simultaneity slice," so that this "all at once" branching, involving, for example, simultaneous branching in Chicago and at a chosen spot on the surface of Neptune, looks to be a serious violation of the relativistic prohibition of "action at a distance." Move 2: The core idea of Special Relativity comes from the Newtonian basis by retaining determinism, while replacing global influences (action at a distance) with relativistic only-local influences. The "world" of Special Relativity is composed of "point events" in a four-dimensional "space–time" based on a causal ordering. The causal future of a point event is restricted to the forward light cone determined by the causal order, and, symmetrically, the causal past is confined to the backward light cone. Finally, the theory of branching space–times is intended to be both indeterministic and relativistic: Many possible histories branching each from each, with each history being a relativistic space–time. Both the indeterminism and the relativity ingredient in the theory of branching space–times are intended as "pre-physics"; for example, no probabilities and no frames of reference—indeed, no numbers. The aim here is to highlight deep features of these theories by suppressing the serious physics and the numbers that they require. In order to avoid both the forbidden notion of absolute simultaneity and also "action at a distance," the branching between two space–times must occur at a point event. It seems clear that a good theory of branching space–times does not at once come to mind.

The primitives of BST are two: (1) *Our World*, abbreviated *OW*, is intended to be a set of "concrete possible point events" (or just "point events") $e \in OW$, and (2) $<$, the "causal order" on *OW*. The causal-ordering relation $e_1 < e_2$, which has both spatio-temporal and modal significance in BST, may be read as either "e_1 is in the settled causal past of e_2" or "e_2 is in the future of possibilities of e_1." It is assumed that $<$ is a strict partial order on *OW* with no maximal elements. Where a *chain* is a subset of *OW* such that two distinct members are always comparable by $<$, an *outcome chain*, O, is a nonempty lower bounded chain, and it is assumed that every outcome chain has an infimum. A *history*, h, is a maximal directed set, where a set is *directed* if it contains an upper bound for each pair of its members. An *initial chain*, I, is nonempty and upper bounded, and it is assumed that every initial chain has a supremum in every history that contains it. Fact: Under the causal ordering, histories are closed toward the past (if two histories overlap at a point event, they share the whole causal past of that point event). H is a set of histories (also called a *proposition*). There is also a postulate saying that given two initial chains and two histories, the order of the respective suprema is preserved as the histories are varied. But far and away the most distinctive assumption is the "prior choice postulate," which says that every contingent outcome has an originating cause in its past:

Prior choice postulate. *Let O be an outcome chain, and let $O \subseteq (h_1 - h_2)$. Then there is a point event e in the causal past of every member of O such that e is maximal in $(h_1 \cap h_2)$.*

These postulates taken together yield "BST," the theory of branching-space times according to Belnap [1].

[1] There are in the literature several candidates for BST theory and several developments thereof. Consult in particular the works of Kowalski and Müller and Placek and Weiner ~~listed in the references~~. Of special note are their contributions to the theory of probabilities in BST, to the theory of counterfactual conditionals in BST, and to the application to certain problems in quantum mechanics.

Please cite this article in press as: N. Belnap, Prolegomenon to norms in branching space–times, Journal of Applied Logic (2011), doi:10.1016/j.jal.2009.09.006

The prior choice postulate leads to a certain concept of causes in branching space–times called *causae causantes* studied in Belnap [6] (henceforth CC). I have to go through some technicalities to get there, so that it will be good to keep in mind the relevant concept of causation at which we are aiming. It comes from the work of Mackie [18]. There Mackie develops an idea which has become well known: the idea of a cause as an "INUS condition," using an acronym for "***i**nsufficient but **n**on-redundant part of an **u**nnecessary but **s**ufficient*" condition (think constituent of a disjunctive normal form). An important difference is that here, in contrast to Mackie, I develop a rigorous theory of causes and effects as certain kinds of *events* (rather than *propositions*) in BST. Here is a bare-bones sketch:

1. $H_{\langle O\rangle}$ is the set of all histories overlapping O; it represents as a "proposition" O's beginning to be, which is the way that outcome chains "occur" in a history.
2. $H_{(e)}$, the set of all histories to which e belongs, is the proposition that e occurs. $h_1 \equiv_e h_2$ is read "h_1 is undivided at e from h_2," and is provably an equivalence relation on $H_{(e)}$. The induced partition of $H_{(e)}$ is written Π_e, with $\Pi_e\langle h\rangle$, for $e \in h$, being the member of Π_e to which h belongs. The idea of Π_e is essential to BST theory. When a set of histories, H, belongs to Π_e, it is proper to understand H as an *elementary possibility at* or *immediate outcome of* e. Of course one case, the uninteresting case, is that $H = H_{(e)}$, making the partition trivial. In the nontrivial case, if H_1, H_2 are distinct members of Π_e, they must be disjoint, which is to say, *inconsistent*. e is then said to be a *choice point*; in this case, if history $h_1 \in H_1$ and history $h_2 \in H_2$, we write $h_1 \perp_e h_2$ to say that the two histories "split" at e: Before and at e, both histories are possible (e belongs to both histories), but at any point event after e, at most one is possible (no point after e belongs to both histories). Splitting at choice points is how *Our World* counts as being a world of "branching histories."
3. "Transition events" are of central importance in the BST theory of causation. Defining a *scattered outcome event* as a set, $\mathbf{O}$, of outcome chains all of which overlap some one history, and defining an *initial event*, $\mathbf{I}$, as a nonempty subset of some one history, a *transition event* is defined as an order pair, written $\mathbf{I} \rightarrowtail \mathbf{O}$, such that $\mathbf{I}$ precedes $\mathbf{O}$ in the causal order in an appropriate sense: Every point event in $\mathbf{I}$ properly precedes some chain in $\mathbf{O}$. The entire theory can be generalized by taking outcomes of transitions as "disjunctive outcome events," represented by a set of pairwise inconsistent scattered outcome events; and it can be specialized by restricting attention to "chain" outcome events.
4. Even more specialized are the crucial *basic transitions* that are at the bottom of BST causal theory. They can be represented as point-event/outcome-chain transitions ($e \rightarrowtail O$), where $inf(O) = e$, or as point-event/outcome-proposition transitions ($e \rightarrowtail \Pi_e\langle h\rangle$), where $e \in h$. The two representations turn out to come to the same thing. In either case, a basic transition is from a point event to an *immediate outcome* of e. Such transitions are the ultimate, uncaused originating causes, the *causae causantes*.
5. But causes of what? The simplest to explain is causation of an outcome chain, O. When a contingent outcome, O, occurs, we should always look for an event in *Our World* where things could have gone either way, either ruling out O altogether, or keeping it possible. In BST terms, that means looking for a point event, e, at which there is a split between some history excluding O entirely and the histories $H_{\langle O\rangle}$ in which O begins to be: $h \perp_e H_{\langle O\rangle}$ for some h. The prior choice postulate guarantees that we shall find at least one such in the past of contingent O.
6. Such an e is a *basic cause-like locus* for O. It turns out, however, that the theory of basic cause-like loci is not easily understood in the absence of an additional postulate, the so-called "no funny business" postulate, which we therefore assume. The postulate simply says that *all* cause-like loci for O lie in the *past* of O. I therefore let $pcl(O)$ be the set of past cause-like loci of O. We might call the resultant theory simply BST + NFB; for the duration of this essay, however, we use plain "BST" as short for "BST + NFB."
7. But which of the many basic outcomes of a past causal locus of e should be considered in defining the idea of a *causa causans* of O? If there were funny business, that question might be difficult, but as it is, since e is in the past of O, we can calculate that the occurrence of O is consistent with (i.e., has a history in common with) *exactly one* basic outcome of e. Let us give it a name: Provided $e < O$, define $\Pi_e\langle O\rangle$, read "the projection of O onto e," as the unique member of Π_e that is consistent with $H_{\langle O\rangle}$.
8. Finally, define $cc(O)$, read the set of *causae causantes* of O, as $\{(e \rightarrowtail \Pi_e\langle O\rangle)\colon\ e \in pcl(O)\}$; that is, as the set each member of which is a nontrivial basic transition with initial e whose outcome is precisely the projection of O onto e.
9. This story generalizes, the most general concept being the set $cc(\mathbf{I} \rightarrowtail \mathbf{O})$ of *causae causantes* of a transition event with initial, $\mathbf{I}$, and disjunctive outcome, $\mathbf{O}$. Details, necessarily omitted here, are all to be found in CC. The critical upshot is that the full set of *causae causantes* of an outcome event or a transition event is *provably* always a full set of INUS conditions of the event in the sense of Mackie [18], where, to repeat, an INUS condition is an *insufficient but non-redundant part of an unnecessary but sufficient condition*. The *causae causantes* as a whole give the complete objective causal story in terms of events that describe the "why" (or perhaps the "how") of the outcome or transition event.
10. The concept of a scattered outcome event, $\mathbf{O}$, is intermediate between that of a chain outcome, O, and a disjunctive outcome, $\mathbf{O}$. The scattering can be either space-like or time-like. Being non-disjunctive, a scattered outcome event, $\mathbf{O}$, or transition event $\mathbf{I} \rightarrowtail \mathbf{O}$, like a chain outcome, O, will have a set of *causae causantes* that are to be taken non-disjunctively, not as a set of INUS conditions, but instead as a set of "INNS" conditions, each of which is an ***i**nsufficient but **n**on-redundant part of a **n**ecessary and **s**ufficient condition*. Some of the most influential literature on causation either confuses INUS and INNS conditions, or, more likely, gives an only informal and partly subjective account of the difference.

Please cite this article in press as: N. Belnap, Prolegomenon to norms in branching space–times, Journal of Applied Logic (2011), doi:10.1016/j.jal.2009.09.006

That may seem complicated, but I think that nothing less complicated will do if what is wanted is an objective and rigorous theory of causation in the events.

3. Agents

FF took the basic locution of the theory of agents and their choices to be [α *stit*: Q], read "α sees to it that Q." There was no separate theory of agents independent of this locution. As a consequence, in FF, based as it was on branching time, and with explicit refusal to import any "mental" concepts, there could be no direct analysis of "the real internal constitution" of an agent. Instead, a nonempty set called *Agent* was postulated, members of which were characterized only indirectly, by reference to, so to speak, what they objectively did or might have done, or might do.

3.1. The representation of the agent

Here I outline a theory, "BSTAC," of how agents and choices fit into branching space–times. Even though the name doesn't so suggest, BSTAC is intended to be limited by the inclusion of a "no funny business" postulate. Because BST takes *OW* to be based on point events instead of moments (which, recall, are instantaneous in time, but maximally spread out in space), something more can be said about the inner structure of an agent.

The simple idea is to adapt the theory of agents to the chief difference between non-relativistic branching *time* and relativistic branching *space–times*, namely, the difference between what "the causal order relation" relates. The critical feature of agents in branching time is that *moments* are super-events taken to be "spatially" rich enough to be "occupied" by more than one agent. That in turn calls for a BT-stit postulate insuring that the simultaneous choices of distinct agents must necessarily be independent. There can be no "influence" of the choice of one agent at four o'clock sharp on the choice of another agent at exactly that same time. To suppose the contrary is to suppose a form of "action at a distance," and to become involved in pointless muddles, perhaps involving some confused idea of "common cause" that is not relevant to (free) choice.[2]

Branching space–times theory, however, postulates *point events* instead of moments. Point events would seem to be so small as to admit the "presence" of at most one agent. Perhaps for modeling some weird scenarios one might wish to play with "double occupancy," but it is a byway that at this point is best left unexplored.[3] For this reason it seems reasonable to begin by representing an agent as a set of point events, the set of point events that may be thought of as possibly occurring in the course of his or her life. (For reasons given at the end of the discussion below of the "External standpoint," it is better to think this way instead of copying branching-time theory by thinking of the point events as "locations" of the agent.) In contrast, in branching time it would make no sense at all to represent an agent as a set of moments! Continuing to use *Agent* as the set of agents, we may therefore begin with a simple assumption.

Postulate 1 *(Representation of the agent). Every agent, $\alpha \in Agent$, is a nonempty set of point events in OW*: $\forall\alpha[\alpha \in Agent \rightarrow \varnothing \neq \alpha \subseteq Our\ World]$.

I hope that it is clear that I intend Postulate 1 as a *representation* of agents, not as an account of what agents "really are."

3.2. Standpoints

To make further sense, it is essential to assume a "standpoint."

3.2.1. External standpoint

We may take ourselves for present purposes to be taking a "godlike" or "scientific" or "external" standpoint entirely outside of *Our World*. In this case all point events are on a par: All are mere possibilities, and we are not entitled to use either tense expressions or differentially applicable modal expressions. From the external standpoint, each member, e, of α represents one of α's *possible* (*point-like*) *life events*. For mere brevity, I usually drop "point like." Perhaps this is an expository mistake, given that I certainly wish to leave theoretical room for extended life events. The shorter phrase should, however, work for this essay. We may say if we like that all point events (hence all of α's possible life events) "are equally real," as is sometimes said with reference to the Everett interpretation of quantum mechanics; but it must be borne in mind that their "reality" is the reality of mere possibilities. From the envisioned standpoint, it makes no sense to distinguish two kinds of point events, the "possible" and the "actual," much less three kinds, including the "impossible."

[2] The postulate works to keep the BT-stit theory out of trouble, but it is after all just wrong to think that choices of agents operating jointly should occur at exactly the same moment. A reason to move from BT to BST was precisely to avoid "trouble" without the wrongness. Section 4.1 calls the BT postulate "ludicrous." It must be observed, however, that BT itself seems to have the resources causally to combine successive choices by one or more agents by applying to BT the ideas of *causae causantes* as transitions—a thought that is not new, but since not having been worked out in detail, has heretofor been expressed only in conversations.

[3] Later we enter Postulate 5, which asserts that "double occupancy" doesn't happen.

Please cite this article in press as: N. Belnap, Prolegomenon to norms in branching space–times, Journal of Applied Logic (2011), doi:10.1016/j.jal.2009.09.006

Nevertheless, the external standpoint welcomes *relational* versions of possibility, actuality, and impossibility of point events. To say that a point event, e_1, is (1) possible or (2) actual or (3) impossible *relative to* e_0, is simply to say that (1) some history contains both e_0 and e_1, or (2) that every history containing e_0 also contains e_1 (that is, that at e_0 it is settled that e_1 occurs), or (3) that no history contains both e_0 and e_1.

In a useful analogy, if (1) we take up a standpoint that is not itself located on the surface of planet Earth, then (1a) to say, relationally, that "*a* is to the north of *b*" is sensible, but (1b) it makes no (ordinary) sense to use "is to the north" as a one-place predicate. If, however, (2) we ourselves are located on the surface of planet Earth, then the context of utterance supplies an otherwise missing second argument to the use of the one-place predicate, "is to the north." As a corollary, if (1) we take up a standpoint external to *Our World*, then (1a) relational "earlier/later," and "space-like related" between pairs of possible life events still make sense, but (1b) one-place "past/present/future/over-there" make no sense. Only when (2) we take up a standpoint in *Our World* do the latter phrases make sense, provided they are understood as context-dependent. (Shades of McTaggart!)

From such a "godlike" or "scientific" standpoint, when $e \in \alpha$, we may say, in a strictly technical voice, that the point event *e is part of* the agent α. One might also say that α *is located at* or *occupies* *e*, but it is much better to say that *e* is one of α's possible life events. The reason is that "location" suggests a position concept, but whereas a point event, *e*, may *have* a position in *OW*, it is not a good idea to say that *e is* a position. After all, a central idea is that two radically incompatible point events, each in a distinct possible future of some chance event, may occupy the same spatio-temporal position. Further, invoking the ordering relation, we may say that when $e_0 < e_1$, either that e_0 is in the causal past of e_1, or that e_1 is in the (causal) future of possibilities of e_0, since these phrases are tenseless and modally constant.

3.2.2. Internal standpoint

We may instead take ourselves to be located at some point event (paradigmatically the one to which "here-now" refers) in *Our World*. This is most easily viewed through the lens of language, and so, idealizing, you should envisage the point event as the *point-event of utterance*, or the *context-of-utterance point event*.[4] In this case, the fundamental "truth" locution should be

$$\mathbb{S}, \mathbb{I}, e_c, e, h \models A, \text{ read as}$$

- *A* is true with respect to structure $\mathbb{S}$, interpretation $\mathbb{I}$, point event of utterance e_c, point event *e*, and history *h*.

Here $\mathbb{S}$ lists the *structure parameters*, especially a domain of quantification and the structural elements *OW* and $<$, but room is left for space–time structural elements such as a "frame of reference." $\mathbb{S}$ may itself be called *the* structure parameter. $\mathbb{I}$ is the *interpretation parameter* whose job is to interpret the atomic pieces of the language, e_c is *the point event of utterance*, *e* is *the point event of evaluation* (an auxiliary parameter moved by tenses), *h* is *the history of evaluation* (moved by the "settled true" construction), and *A* is the sentence to which parameter-dependent truth is attached. I skip the details, except to emphasize that the causal future tense requires that, at a minimum, truth be relativized to a point event, *e*, and a history, *h*, that contains *e*. How this should go can partly be inferred from Chapter 8 of FF, *mutatis mutandis*. See also the use of "double time references" in Section 6 below; and for a general account of parameterized truth, see Belnap [8].

3.3. Agents and world lines

Given that we represent an agent, α, as a set of point events to be thought of as α's possible life events, what constraints make sense? In the beginning it seems best, since easiest, to think of the life of an agent in a particular history as a portion of a "world line." In other words, for each history, *h*, if we take the set of those of α's possible life events that occur in *h*, they form a linear order, stretching from α's birth to α's death in *h*, each causally comparable with each. (Whether it counts as a fault or not, I confess little or no interest in imaginative speculations concerning a "personal split" within a single, consistent course of events; within, that is, a single space–time.)

Postulate 2 *(Agents and world lines). The portion of the life of an agent in a particular history is a chain of point events*: $\forall\alpha\forall h$ [$(\alpha \in \mathit{Agent} \,\&\, h \in \mathit{Hist}) \rightarrow (\alpha \cap h)$ *is a chain in Our World*]. *If* $(\alpha \cap h) = \varnothing$, *we may say that* α *"doesn't exist" in h. For each agent,* α, *and history, h, we call the set* $(\alpha \cap h)$, *provided it is nonempty, the* life history *of* α *in h, and therefore a* possible life history *of* α.

Postulate 2 is at this stage of sophistication a natural and helpful postulate, but it is as well to recognize that in a longer run it is likely to be a drastic simplification of the notion of life history. I hope it is needless to say that I am claiming for this postulate only that its simplicity makes it a good beginning; it may well turn out to be useful to represent an agent in a single history as a cloud of point events rather than as a chain. Our chain representation rules out, for instance, that distinct bodily point events count as distinct parts of a life history, since many such points are space-like related and accordingly

[4] The phraseology is barbarous, but I have not been able to think of a safe improvement. Mere "point," for instance, would wrongly suggest a position in space–time, a position subject to being occupied by any one of a variety of pair-wise inconsistent concrete possible point events. I do in fact reserve "point" for this positional use.

Please cite this article in press as: N. Belnap, Prolegomenon to norms in branching space-times, Journal of Applied Logic (2011), doi:10.1016/j.jal.2009.09.006

break linear order. Let us, however, leave bodily parts for another day. The thought is that Proposition 2 supplies a useful and manageable first approximation.

It follows at once that α is a tree. That is, α is partially ordered by $<$ (since α is a subset of *OW*), and α satisfies no backward branching (that is, it's false that for some $e_0, e_1, e_2 \in \alpha$, $e_0, e_1 < e_2$, but neither $e_0 \leqslant e_1$ nor $e_1 \leqslant e_0$). For suppose $e_0, e_1, e_2 \in \alpha$, and $e_0 < e_2$ and $e_1 < e_2$. e_2 must belong to some history, say h, and therefore $e_0, e_1 \in h$ by downward closure of histories. Hence $e_0, e_1 \in (\alpha \cap h)$, and since $(\alpha \cap h)$ is a chain, either $e_0 \leqslant e_1$ or $e_1 \leqslant e_0$.

Next, following out an intuition that goes back at least to [16], we assume that each agent has a unique birth event, and in each history a unique death event.

Postulate 3 *(Birth and death of agents). If α is an agent, there is a historically (or modally) unique point event that is least among all members of α. This we call the* birth *of α. Furthermore, for each history, h, such that $\alpha \cap h \neq \emptyset$, the life history of α in h has a maximum (a unique greatest member). This we call the* death *of α in h.*

On this construction, the death of α in h is a "life event" of α; not only does that seem technically harmless, but informally both birth and death are often listed as events in the life of an agent.

The asymmetry in the representation of birth and death does not come from casual English, which, given an event, e, in the actual life of, say, Alexander Hamilton, is equally happy with alternative possible futures vs. alternative possible pasts of e. The former sensibly underwrites alternative possible death events, d_1 and d_2, of Alexander Hamilton, but BSTAC properly forbids the latter, which would underwrite (distinct) alternative possible birth events, b_1 and b_2. (Evidently the least member of a tree must, if it exists, be unique.)

Is it a strength or a weakness of BSTAC that it makes impossible that Hamilton was born earlier or later than he in fact was? There is no doubt that in BSTAC, if we take "Alexander Hamilton" to name a particular agent, then Hamilton's birth event is modally fixed. On the other hand, nothing so far laid down prohibits the development of a disjunctive concept of an agent with different birth-events for each component. How that would go I do not know. I suspect that opening up the theory in this way would lead to loose talk; but that is only a suspicion.

I also assume that the life history of each agent in each history is continuous:

Postulate 4 *(Continuity of life histories).*

(1) *The causal ordering relation is dense in each life history of each agent*: *For $e_0, e_2 \in (\alpha \cap h)$, if $e_0 < e_2$, then there is a point event e_1 in $(\alpha \cap h)$ such that $e_0 < e_1$ and $e_1 < e_2$.*
(2) *For each agent, α, and history, h, if E is a subset of the life history of α in h, then the infimum of E belongs to that life history, and so does the supremum of E relative to h.* (*No jumps or gaps.*)

This is perhaps more "metaphysics" than we need, but it makes our representation of agents easier on the imagination.

It is easy to become confused when speaking of persons (agents) against a background of indeterminism such as is coded into branching space–times. It helps to remain aware that an agent is represented at three levels.

3.3.1. Summary of three levels

1. A *possible life event*. We may say that e is one of α's *possible life events* iff $e \in \alpha$ (where e is a point event in *OW* and α is an agent in *Agent*).
2. A *possible life history*. We may say that a subset E of *OW* is the *possible life history* of α in h iff $E = \alpha \cap h$. By natural quantifications, one may speak of "α's possible life histories in *Our World*," and "the possible life histories of agents in h."
3. An *agent* (or a person). α is an *agent* iff $\alpha \in Agent$, and hence $\alpha \subseteq OW$. On this representation, each alternate life history is a literal part of the agent. In slightly looser language, we might say that alternate possibilities are ingredient in the very nature of the agent. This is to be understood in a sense consistent with the doctrine that some *causae causantes* of a life event of agent α may be extrinsic to α, as we see in Section 3.4 below.

All of this is from the external point of view. From that standpoint, "possible" draws no distinction. For that reason, we can, if we like, omit "possible" from "possible life event" and "possible life history" without doing much harm.

It is evident that the critical difference between (1) a life history of an agent and (2) an agent is this: Only an agent contains not just choice points, but also the multiple jointly inconsistent but individually possible outcomes of choices sharing a given choice point as initial. Put from a certain internal standpoint, thus is enabled "She chose to walk over the bridge, but if instead she had chosen to run, she would have arrived in time." No single life history can ground such a story.

When we are anchored to some particular context of utterance, e_c, we may speak of "the actual past life history of α" and "the future (portion of the) possible life histories of α," etc. "Actual" and "past" and "future," when not explicitly relational, require anchoring to a context of utterance, and *you may not use them if your standpoint is external*—unless you are prepared to meet the "no thin red line" arguments of Chapter 6 of FF.

Please cite this article in press as: N. Belnap, Prolegomenon to norms in branching space–times, Journal of Applied Logic (2011), doi:10.1016/j.jal.2009.09.006

3.4. Branching of life histories

Next I explore branching of the life histories of an agent, α.

I've idealized a possible life of an agent as a kind of spatio-temporal "worm" in a history. This representation is familiar since each individual history is supposed to be a space–time, and for a very long time philosophers—and often scientists—have played with the worm picture. We then saw that the representation of each possible life history as a chain or worm implies that the agent as a whole has the shape of a tree (no downward branching). It is tempting to leap to the conclusion that each branching of the life histories of an agent, α, represents a decision point for α, but that would be wrong. BST theory is more subtle than that. What is true according to the theory is the "prior choice postulate," according to which if some life event, e_1, for an agent, α, occurs in one history, say h_1, but not in another history, say h_2, then there is in the past of e_1 a (possibly metaphorical) "choice point," e_0, which is a branch point in the sense of being maximal in the intersection $h_1 \cap h_2$ of the two histories. There is no warrant, however, for believing that e_0 is one of α's life events, or indeed is a life event of any agent whatsoever. The indeterminism might be resident instead in some thoroughly non-agentive event such as a fair toss of a coin or, perhaps, a choice point between spin-up and spin-down for some electron.

It remains true that a representation of an agent, α, must look a lot like a tree, but α's tree will in general have two quite different kinds of branching. Which kind will depend on whether the indeterminism represented by the branching is a choice in the life of α, or whether instead the locus of the indeterminism is foreign to α. In both kinds of branching, there is a single past-pointing worm-like representation of the past life of α up to the branching, and an entire assemblage of distinct worm-like representations of the possible future-lives of α subsequent to the branching, severally distributed among the histories in which the life of α continues.

More specifically, let α exist in both of two histories, h_1 and h_2, but with different deaths. In other words, let $(\alpha \cap h_1) \neq \varnothing$ and $(\alpha \cap h_2) \neq \varnothing$ and $(\alpha \cap h_1) \neq (\alpha \cap h_2)$. Since histories are closed backward, the intersection $(\alpha \cap h_1 \cap h_2)$—call it E_0—must be a nonempty backward-pointing tail running from immediately below the branching clear back to the birth of α. Let E_1 be the chain $(\alpha \cap h_1 - h_2)$, and let E_2 be the chain $(\alpha \cap h_2 - h_1)$. $L_1 = (E_0 \cup E_1)$ and $L_2 = (E_0 \cup E_2)$, being respectively $(\alpha \cap h_1)$ and $(\alpha \cap h_2)$, must each be a chain running from birth to death. By continuity and Dedekind, either (1) E_0 has a maximum, e_0, or (2) each of E_1 and E_2 has a minimum, respectively e_1 and e_2. In case (1), one can show that not only is e_0 a choice point for the chains L_1 and L_2, but e_0 can also be shown to be a choice point for h_1 and h_2, that is, a point event maximal in $h_1 \cap h_2$. In this case, since $e_0 \in \alpha$, we are entitled to attribute the split between h_1 and h_2 to the agency of α, at least in part (there may be other maxima in $(h_1 \cap h_2)$). In case (2), however, the minimum e_1 of E_1, which is in $h_1 - h_2$, must be a supremum of $E_0 \cap h_1$, and symmetrically the minimum e_2 of E_2, which is in $h_2 - h_1$, must be a supremum of $E_0 \cap h_2$. So E_0 must have two (distinct) history-dependent suprema, one in h_1 only, and the other in h_2 only. These two suprema, being as close as they are to E_0, that is, each with nothing in between, must themselves be vanishingly close to each other. Although OW is topologically a $T1$ space, e_1 and e_2 testify that OW is not a $T2$ space; that is, OW is non-Hausdorff and thus explicitly not a manifold. If one had a doctrine of "space–time positions," e_1 and e_2 would have to occupy the same space–time position, differing only modally by being in different histories.[5]

This representation of agency in branching space–times presupposes continuity of each space–time. How best to represent agents-in-branching-space–times discretely, as is perhaps required for computer applications, is open for research.[6] But just as computer representations of real arithmetic must in some way answer to "real" arithmetic of the reals, so any such discrete representation must in the end answer to *Our World* as a BSTAC.

4. Choices

At this point we have a theory of *pcls* and *causae causantes* from Section 2 and a theory of agents from Section 3. (Recall that a "*pcl*" or "*past causal locus*" of O, is a point event in the past of O at which one possible immediate future of e is consistent with the proposition that O occurs, whereas the rest of the possible immediate futures of e are inconsistent with the occurrence of O.) These analyses leave us little option for an analysis of a *choice point* for an agent, α, relevant for an outcome, O; namely,

Definition 1 *(Choice point, favoring).*

(1) e is a *choice point for* agent α *relevant for* outcome O iff e is a life event of agent α that is a past causal locus of O; that is, $e \in (\alpha \cap pcl(O))$.
(2) Provided e is a choice point for α relevant for O, $(e \rightarrowtail \Pi\langle O\rangle)$ is *a choice by* α that *favors* O.
(3) A life history of α *is relevant for* O if some life event of that life history is a past causal locus of O.

[5] Please do not infer that OW is a weird space–time. It is not, of course, a space–time at all, neither branching nor non-branching; it is, in intent, an *assemblage* of space–times that branch one from another, the branching being modal rather than spatio-temporal. Earman [12] endorses "ensemble" branching in contrast to "individual" branching. Since Earman, however, loads his phrase with a special meaning, it seems best to avoid "ensemble" in favor of "assemblage."

[6] I say "perhaps" because the continuity of each space–time by no means demands that every model admit an infinity of agents or of *causae causantes*, much less a continuum of such.

Please cite this article in press as: N. Belnap, Prolegomenon to norms in branching space-times, Journal of Applied Logic (2011), doi:10.1016/j.jal.2009.09.006

All this is spoken from an external standpoint. Observe that the choice is a *transition* event, that is, an ordered pair satisfying certain conditions. Why a transition event? Well, an *initial* event does not do justice to the idea of choice, because no outcome of the choosing is specified. If you are on the verge of deciding to turn left, you are also (say) on the verge of deciding to turn right, so being "on the verge" cannot count as the choice. So also an *outcome* event, you having decided to turn left, although a better candidate (since its infimum must be the initial of the choice), does not on the face of it represent the particular, concrete indecision between (say) turning left and turning right that is certainly part of the choice. Reification of a choice ought to represent both the undecided initial and the result of the decision; that's why it is represented as an ordered pair of initial and outcome. In other words, at any time up to and including (say) t, there is more than one possibility for the future, and at any time after t the choice has already been made; so when was the choice? You need both a "before-and-up-to t" with no-choice-yet, and an "after t" with choice-has-been made; that is, you need a transition. A lot less ink would be spilled on the concept of choice if this were recognized as a truism.

In Definition 1(2), the idea of "favoring" is a notion of "partial causation." A transition "favors" an outcome if it keeps the outcome possible at least for a while; that is, although there is no thought that the transition guarantees the occurrence of the outcome, O, it does not render it henceforth impossible—as does every other immediate outcome of e.

4.1. Stit

The workhorse locution "$[\alpha\ stit{:}\ \ Q]$" in FF, where α names an agent and Q is a sentence, signifies that the choice of the agent, α, guarantees the truth of Q (Positive condition) in circumstances in which Q is not already settled true (Negative condition). If we are speaking of a single choice, it needs observing that it must be seldom indeed that a single basic transition, $(e \rightarrowtail \Pi_e\langle h\rangle)$, can guarantee the truth of Q, unless Q reports, in effect, an immediate outcome of e, a case that I first consider.

I rely on the fundamental truth locution, $\mathbb{S}, \mathbb{I}, e_c, e, h \models A$, pertinent to the "internal standpoint" of Section 3.2. The only difference from FF is the move from the moments of branching time to the point events of BST. The account of *stit* is structurally the same.

> $\mathbb{S}, \mathbb{I}, e_c, e, h_1 \models [\alpha\ stit{:}\ \ Q]$ iff $e \in \alpha$ and $\mathbb{S}, \mathbb{I}, e_c, e, h_2 \models Q$ for all h_2 such that $h_1 \equiv_e h_2$ in the sense of Section 2(2) (Positive condition), and $\mathbb{S}, \mathbb{I}, e_c, e, h_2 \not\models Q$ for some h_2 such that $e \in h_2$ (Negative condition).

This is the *dstit* of FF and [13], which looks better in branching time than in branching space–times, since in the former case, originating causes must be linearly ordered. When, however, we consider many *causae causantes* based on many past causal loci in *Our World*, it is forced upon us that the various *pcls* are likely to be spread out, some pairs having a space-like relation, and other pairs a causal relation. It's not easy to see the point of collapsing this wealth of structure into the *stit* framework, which was predicated on a *single* moment as initial, even when considering multiple agents. (For a seeing-to-it by a group Γ of agents, FF imposed the artificial and even ludicrous requirement that the choices of the agents in Γ be simultaneous so that the causal independence of the various choices could be guaranteed.)

What we can do in a straightforward manner in BSTAC theory is to catch the notion that an outcome O is entirely due to choices of an agent, α. That gives us a thought connection with *stit*. The practice has been to differentiate *stit* concepts by means of a prefixed letter, as in *astit*, *dstit*, and so on. In analogy, I use "*pstit*" for a *stit* concept based on past causal loci, so that $[\alpha\ pstit\ O]$ might be read as "α sees to it that Q in virtue of a system of prior choices by α."

Definition 2 *(pstit in BSTAC).* $[\alpha\ pstit\ O]$ iff $\varnothing \neq pcl(O) \subseteq \alpha$.

Since every *causa causans* of O must originate in a past causal locus of O, the definition says that every *causa causans* of O is identical to a choice by α that favors O, so that the choices by α give a complete causal account of the occurrence of O. The choices by α taken together give a sufficient condition for O, and each choice separately is (in general) insufficient but necessary and (most critically) is irredundant; so that we have an INNS condition.

It is necessary to observe that Definition 2 defines an "external" phrase, whereas the aim of FF was to characterize a phrase that could serve as a context-dependent part of a language to be spoken by those taking up a standpoint in an indeterministic world. The tasks are not equivalent; starting from Definition 2, to determine a happy path to such a characterization is by no means automatic. Observe, for instance, that the present story about seeing-to-it-that in branching space–times is more narrowly focused than the companion story of FF about *stit* in branching time: The Q in $[\alpha\ stit{:}\ \ Q]$ took the place of an arbitrary sentence, which could express an arbitrary proposition, but Definition 2, in contrast, explains "seeing to it that in virtue of a system of prior choices" only for outcome events, and indeed only for outcome chains, the easiest case of all, and so derivatively only for rather simple outcome propositions.

Still, the system of *causae causantes* is in certain dimensions much less limiting than that of FF, and that not only because of the move from branching time to branching space–times. I can illustrate what I mean while staying entirely in the branching-time framework. The scheme of FF is defective to the extent that there is no way to represent an outcome as caused by a succession of contributory causes. There is no way to say that Mary saw to it that the dishes were washed by carrying out a sequence of choices: First she chose to do the silverware at moment m_1, then the plates at m_2, and, finally, the serving dishes at m_3. The scheme of FF permits only the last to count, the others being redundant just by being causally

Please cite this article in press as: N. Belnap, Prolegomenon to norms in branching space–times, Journal of Applied Logic (2011), doi:10.1016/j.jal.2009.09.006

earlier and hence already causally implied by the occurrence of that final outcome. The thought is not that in general choosing to wash the serving dishes implies having previously chosen to wash the plates. Instead, since histories are closed downward, there is a bare geometric fact: That the final choice point, m_3, occurs in a history, h, implies as a "historically necessary condition" not only that m_2 (which was postulated as earlier that m_3) also occurs in h, but furthermore that an outcome chain, O, representing the outcome of the plates choice at m_2, also occurs in h. That's because there (provably) is a *unique* immediate outcome of m_2 that favors m_3.

The present analysis overcomes this severe defect by not looking at mere outcomes (*silverware washed, plates washed,* and *serving dishes washed*), but by instead looking at the basic transitions, ($m_1 \rightarrowtail$ *silverware-washed*), ($m_2 \rightarrowtail$ *plates-washed*), and ($m_3 \rightarrowtail$ *serving-dishes-washed*). It is those transitions which, when properly understood, are each an *irredundant* necessary condition of her washing the dishes.[7] Those transitions, taken together, form a set of INNS conditions (INNS rather than INUS because, in this example, not disjunctive).[8]

All is not, however, plain sailing. It remains natural to ask *when* Mary washed the dishes. If the question is asked rhetorically as part of an effort to cast doubt on BSTAC theory, the defensive response is to remind the questioner that we are dealing with transitions, and that we have all known from the beginning that transitions, to use an apt expression coined by a famous Philosopher, have no "simple location." To suppose otherwise is to commit Whitehead's "fallacy of simple location." One might suggest that it is unlikely to be fruitful to inquire into, for example, "the time of a killing." That negative-minded suggestion may well have, however, a variety of interesting counters. I leave the matter in this open state.

5. Many agents

One agent is hardly enough to populate *Our World*: The theory must provide for many agents. This is a complicated matter, needing additional research. Here I say just a few easy things.

Let us think of the joint representation of two or more agents. The thought comes to mind that two agents might share a life event. I think, however, that no good will come of tracing out the consequences of this thought. It is just too weird, too spooky. Not that it might not happen; it's just that if it does happen, present thinking is not likely to help clarify the matter. I therefore enter the following.

Postulate 5 *(No agent overlap). Agents never overlap*: $\forall\alpha_1\forall\alpha_2[(\alpha_1, \alpha_2 \in \mathit{Agent} \,\&\, \alpha_1 \neq \alpha_2) \rightarrow (\alpha_1 \cap \alpha_2) = \varnothing]$.

The postulate says, in effect, that the"world lines" of two agents cannot intersect. It is something like "no two material objects in the same place at the same time," and perhaps that observation is enough to raise the comfort level of the postulate.

Agent is somewhat like an "absolute concept" in the sense of Bressan [10]: If you are given a life event, e, there is a unique member of *Agent* of which e is a part. In other words, if $(\alpha_1 \cap \alpha_2) \neq \varnothing$ then $\alpha_1 = \alpha_2$.

It is easy to say, from the external standpoint, that a group of agents is entirely responsible for O:

Definition 3 *(Group stit).* Provided $\Gamma \subseteq \mathit{Agent}$, $[\Gamma \ \mathit{stit} \ O]$ iff $\varnothing \neq \mathit{pcl}(O) \subseteq \bigcup\Gamma$.

Thus one can see that the apparatus of *pcls* and *causae causantes* makes it effortless to move from saying that the occurrence of O is entirely the doing of a single agent, α, to saying that the set of agents, Γ, is fully responsible for the occurrence of O.[9] Even so, although it is easy to say, one would have to choose O carefully in order to have an example that rang true. Idealization will have to play a role in applying any of these *stit* concepts. There are a variety of such examples to be found in FF, and although worked out there in branching time, given enough leeway for idealization, there should be no problem in transporting them to BSTAC. In contrast, the analogous *favoring* concepts should have wide application, the point being that virtually all that we "do" requires the "cooperation" of other agents, and of nature as well. (Please put the blame for this on *Our World*, not on me.)

5.1. Message sending and receiving

To put in relief one aspect of group responsibility, one would expect that there be *communication* among members of the group. What can the theory of branching space–times say about communication, while still holding to its policy of foregoing mental concepts, which are inevitably loose? What, that is, can we say about the causal underpinnings of sending and receiving and acting upon messages?

[7] To be fit for service in INNS and INUS conditions, it unexpectedly turns out that a transition should be taken to "occur" just when if the initial occurs, then so does the outcome—with the "if–then" understood as a *material* (truth-functional) implication!

[8] FF aimed at something like this with the "transition *stit*" of Chapter 8G.5, but that section entirely missed the mark, and should be consigned to the proverbial flames.

[9] Definition 3 would need to be amended if one wished to prohibit "free riders" and the like. Chapter 10 of FF should be consulted.

Please cite this article in press as: N. Belnap, Prolegomenon to norms in branching space-times, Journal of Applied Logic (2011), doi:10.1016/j.jal.2009.09.006

It seems right that a theory of communication should be developed against an indeterministic causal background, and that for several reasons: (1) There is the question of "noise," construed as randomly generated. (2) Sending a message ought to involve *choice*. I hope you agree that in the present context automated "message-sending" systems are beside the point. (3) Choice is also involved in *receiving* a message. I have in mind what J. Thomson calls "uptake." At the very least, the receiver has to choose to *prepare* for the reception. (4) And if the message is to be "effective," it should contribute causally to the receiver's taking some action. The theory of probabilities in branching space–times (as in Placek [24] and, especially, Müller [20] and Weiner and Belnap [29]) would doubtless play an important role in the theory of message-passing. Stable, detailed considerations on this topic remain to be developed.

6. Norms

This essay was initially conceived as being concerned with norms in branching space–times, and that remains the principal aim of the research of which it is the product. There turned out to be, however, so much required in the way of preliminary conceptual analysis that I have had to shrink what I have to say explicitly about norms to negligible proportions. This accounts for my labeling this essay as a "prolegomenon."

BSTACN (BSTAC with norms) postulates *generated* norms, that is, norms that are generated by a particular localized act of an agent. Wansing [28] has studied such norms in the context of branching time, with the *stit* apparatus, emphasizing the dual roles of the giver and receiver of obligations as agents, etc. Also FF has several chapters that consider norms in branching time, and Müller [21] makes yet further advances. Here my remarks find their basis in "double time references" as described in FF (see the index of FF) and Belnap [2] (DTR). The norm-generating act might be the making of a promise, the laying on of an obligation, the issuing of an invitation, etc. Say that the norm has been *issued*. Such an act need not be a *speech* act, but it is technically convenient to suppose that each norm is issued as if by the use of a declarative core in direct speech. For example, at a point-event, e_0, on Monday, Jack promises Sarah with the following words. *I promise you as follows*:

(∗∗) If it rains in Chicago on Tuesday, I will see to it that you receive $5 on Wednesday.

(∗∗) is called the *declarative core* of the promise-event occurring at e_0. The promise has an interesting semantics that *requires* indeterministic BSTACN as background. It seems essential that at the point event e_0, (∗∗) is neither settled true nor settled false. (Don't say that (∗∗) is neither true nor false; if you do, you will have a hard time avoiding confusion.) To say that the promise made at e_0 on Monday has been carried out at a certain later point event, e_1, on Tuesday means (roughly) that the conditional (∗∗) is true with respect to every point-event-history pair e_0/h, for every history h to which (*not* just e_0 but instead) e_1 belongs, and relative to speaker = Jack, auditor = Sarah, and point-event of utterance = e_0. In other words, at e_1 it is settled true that (∗∗) was (plain) true (*not* settled true) at the point, e_0, at which the promise was issued. This recipe invokes what FF and DTR call a "double time reference," and which MacFarlane [17] describes in terms of branching time by saying that (∗∗) is "true" with respect to moment of utterance e_0 and moment of assessment, e_1. MacFarlane's analysis deserves endorsement; but I beg leave to doubt the wisdom of his terminology, which uses "true" for what is at bottom a modal notion involving quantification over histories.

A description of the norm involved requires more than just double time references. A satisfying—though hardly unique—representation of the normative content of the promising should turn out to be an assignment of "strategies," in the austere sense of Chapters 13 and 14 of FF, to agents in a world of branching space–times. A strategy for α rooted in e_0 will be a function, s, with domain a subset of α and limited to life events at or (causally) after e_0. For e_1 in its domain, $s(e_1)$ will tell α which immediate outcomes of e_1 are permitted by the strategy, so that $s(e_1) \subseteq \Pi_{e_1}$.

The strategic content of the example might be characterized informally as follows. Let the promissor be α_1 and the promisee be α_2, with respective strategies s_1 and s_2, and pick any possible α_1-life-event, say e_1, in the future of possibilities of the promise-event, e_0. If at e_1 it is settled that (∗∗) was true at e_0 (note: Settled-Was, *not* Was-Settled!), the norm has been satisfied, and there is nothing more for α_1 to do. This can be represented by $s_1(e_2) = \Pi_{e_2}$ (anything goes) for all possible life events e_2 for α_1 that are subsequent to e_1. If, however, it is settled at e_1 that (∗∗) was false at e_0, then the norm has been violated, and appropriate sanctions are due. Perhaps $s_1(e_2)$, for $e_1 < e_2$, limits choices by α_1 at e_2 in a punishing sort of way (no more chocolate-chip cookies). Or perhaps the strategy s_2 for α_2 kicks in, permitting or even requiring certain choices by α_2 that favor sanctioning α_1. There are of course many possibilities for cashing out a sanction on α_1.

Aside from the cases in which at e_1 (∗∗) is either settled true or settled false, there remains the crucial case when at e_1 (∗∗) is simply not settled at all; that is, whether at e_1 (∗∗) was true at e_0 is open (historically contingent). Then the strategy, s_1, for the promissor at e_1 requires α_1 to *do* something, at the very least to make a choice among the options Π_{e_1} that keeps the future satisfaction of the promise possible, and, furthermore, if some members of Π_{e_1} guarantee satisfaction of the promise, then to choose one of those. (Without belaboring the obvious, permit me to note that the requirement that α_1 do something is a strikingly objective explanation of the mere metaphor, "world-to-words fit"—nor am I aware of any other comparably objective account.) As indicated in DTR, other sorts of norms call for related descriptions in branching space–times.

Please cite this article in press as: N. Belnap, Prolegomenon to norms in branching space–times, Journal of Applied Logic (2011), doi:10.1016/j.jal.2009.09.006

7. Summary

Here, briefly, is where we have been.

The underlying rationale for this essay given in Section 1 is that norms require agents and indeterminism, that generated norms reference choice events that influence only the causal future, and that an account of norms should give due consideration to space-like-related agent choices. A "branching space–times" (BST) structure that is indeterministic, relativistic, and has room for agents seems to be required. Although not obvious in this largely informal essay, *complete mathematical rigor* is imposed as a constraint on BST theory.

From previous work on BST, I emphasize in Section 2 the ontology of point events and a causal ordering that is a common generalization of the orderings of special relativity and of "branching times." BST theory is a story about "events." A "history" is defined as, in effect, a maximal consistent set of (not propositions but) point events; histories carry the modal burden of BST in strict analogy to "possible worlds" theory (but it's *only* an analogy). Defined kinds of events include initial events, outcome events, transition events, and among the latter, basic "immediate" transitions, which can serve as originating causes, or *causae causantes*. An initial of a *causa causans* of an outcome event is called a "past causal locus," and the set of *causae causantes* of an outcome event provably form a set of INUS conditions in Mackie's sense: Each is an insufficient but non-redundant part of an unnecessary but sufficient condition of the occurrence of the outcome.

An ontology of agents is forthcoming in Section 3, yielding the theory "BSTAC" (branching space–times with agents and choices). An agent is represented in BSTAC as a kind of tree, each branch of which is a continuous chain of point events, from birth to death, in some history. Given an agent, such a chain is called a "life history," and is conceived as consisting in point-like "life events."

As indicated in Section 4, of principal interest with respect to some outcome are those of its past causal loci that are part of the agent. A choice made by an agent might or might not "favor" a certain outcome. When all past causal loci of an outcome lie within the agent, since the agent contains a complete set of INUS conditions for the outcome, we say that the agent "sees to it that" [*stit*] the outcome occurs, thus making contact with earlier work on agents in branching time.

There is in Section 5 a brief discussion of multiple agents, including the observation that distinct agents never overlap in BSTAC, and that generalizing "stit" to a group of agents is straightforward. This idea of a "group stit" is, however, weak. It could be somewhat strengthened (without violating the "no loose mental talk" constraint on BSTAC theory) by the notion of "message passing," which can have an entirely objective account in BSTAC.

Section 6 hints at a theory of norms generated by choices of an agent, for example, the norm that is generated by the making of a promise. The essay advances the thesis that understanding something like a promise requires the ability of BSTAC to handle correctly the "double time reference" involved in saying that a promise made at an earlier point event has (or hasn't) been carried out at a later point event. Finally, the essay suggests that in describing a norm generated by, for example, a promise, one wants to combine double time references with the austere theory of "strategies" of FF, Chapters 13 and 14.

Uncited references

[15] [19] [23] [25] [27]

References

[1] N. Belnap, Branching space–time, Synthese 92 (1992) 385–434. See Belnap [4].
[2] N. Belnap, Double time references: Speech-act reports as modalities in an indeterminist setting, in: F. Wolter, H. Wansing, M. de Rijke, M. Zakharyaschev (Eds.), Advances in Modal Logic, vol. 3, World Scientific Co. Pte. Ltd., Singapore, 2002, pp. 37–58. A preprint of this essay may be obtained from http://www.pitt.edu/~belnap.
[3] N. Belnap, EPR-like 'funny business' in the theory of branching space–times, in Placek and Butterfield [26], pp. 293–315. A preprint of this essay may be obtained from http://philsci-archive.pitt.edu, 2002.
[4] N. Belnap, Branching space–time, postprint, January 2003. This is a postprint of Belnap [1] that includes a number of additional explanations and a little re-structuring. It may be obtained from http://philsci-archive.pitt.edu, 2003.
[5] N. Belnap, No-common-cause EPR-like funny business in branching space–times, Philosophical Studies 114 (2003) 199–221. A non-quotable preprint may be obtained from http://philsci-archive.pitt.edu.
[6] N. Belnap, Agents and agency in branching space–times, in: D. Vanderveken (Ed.), Logic, Thought and Action, Springer, Dordrecht, 2005, pp. 291–313.
[7] N. Belnap, A theory of causation: *Causae causantes* (originating causes) as inus conditions in branching space–times, British Journal for the Philosophy of Science 56 (2005) 221–253. A preprint of this essay may be obtained from http://philsci-archive.pitt.edu.
[8] N. Belnap, An indeterminist view of the parameters of truth, in: T. Müller (Ed.), Philosophie der Zeit. Neue analytische Ansätze, Klostermann, Frankfurt a.M., 2007, pp. 87–113.
[9] N. Belnap, M. Perloff, M. Xu, Facing the Future: Agents and Choices in Our Indeterminist World, Oxford University Press, Oxford, 2001.
[10] A. Bressan, A General Interpreted Modal Calculus, Yale University Press, New Haven, 1972.
[11] D. Davidson, G. Harman (Eds.), Semantics of Natural Language, D. Reidel Publishing Company, Dordrecht, 1972.
[12] J. Earman, Pruning some branches from "branching spacetimes", in: D. Dieks (Ed.), The Ontology of Spacetime II, Elsevier, Amsterdam, 2008, pp. 187–205, Chapter 10.
[13] J.F. Horty, Agency and Deontic Logic, Oxford University Press, Oxford, 2001.
[14] R. Kane, The Significance of Free Will, Oxford University Press, Oxford, 1998.
[15] T. Kowalski, T. Placek, Outcomes in branching space–time and GHZ–Bell theorems, British Journal for the Philosophy of Science 50 (1999) 349–375.
[16] S. Kripke, Naming and necessity, in Davidson and Harman [11], 1972, pp. 253–355.

Please cite this article in press as: N. Belnap, Prolegomenon to norms in branching space-times, Journal of Applied Logic (2011), doi:10.1016/j.jal.2009.09.006

[17] J. MacFarlane, Future contingents and relative truth, The Philosophical Quarterly 53 (2003) 321–336.
[18] J.L. Mackie, The Cement of the Universe, Oxford University Press, Oxford, 1974.
[19] T. Müller, Branching space–time, modal logic, and the counterfactual conditional, in Placek and Butterfield [26], 2002, pp. 273–291.
[20] T. Müller, Probability theory and causation: A branching space–times analysis, British Journal for the Philosophy of Science 56 (2005) 487–520.
[21] T. Müller, Living up to one's commitments: Agency, strategies and trust, Journal of Applied Logic 6 (2008) 251–266.
[22] T. Müller, N. Belnap, K. Kishida, Funny business in branching space–times: infinite modal correlations, Synthese 164 (2008) 141–159.
[23] T. Placek, Stochastic outcomes in branching space–time: Analysis of Bell's theorem, British Journal for the Philosophy of Science 51 (2000) 445–475.
[24] T. Placek, Partial indeterminism is enough: A branching analysis of Bell-type inequalities, in Placek and Butterfield [26], 2002, pp. 317–342.
[25] T. Placek, On Belnap's branching space–times, in: J. Hintikka, T. Czarnecki, K. Kijania-Placek, T. Placek, A. Rojszczak (Eds.), Philosophy and Logic. In Search of the Polish Tradition, Kluwer Academic Press, Dordrecht, 2003, pp. 77–92.
[26] T. Placek, J. Butterfield (Eds.), Non-Locality and Modality, Kluwer Academic Press, Dordrecht, 2002.
[27] T. Placek, T. Müller, Counterfactuals and historical possibility, Synthese 154 (2007) 173–197.
[28] H. Wansing, A reduction of doxastic logic to action logic, Erkenntnis 53 (2000) 267–283.
[29] M. Weiner, N. Belnap, How causal probabilities might fit into our objectively indeterministic world, Synthese 149 (2006) 1–36.

Please cite this article in press as: N. Belnap, Prolegomenon to norms in branching space–times, Journal of Applied Logic (2011), doi:10.1016/j.jal.2009.09.006